Data Acquisition and Trigger Development for RHIC and PHENIX Upgrades

C.-Y. Chi, J. Haggerty, and J.L. Nagle
for the PHENIX DAQ and Trigger Groups*

April 25, 2008

Abstract

In response to a request from the PHENIX Operations Manager,we have studied the implications of increased luminosity at RHIC and the planned PHENIX upgrade detectors on the performance of the data acquisition and trigger systems.

New PHENIX detectors and increased RHIC luminosity from beam cooling will limit the event rate that the PHENIX experiment will be able to record as low as 2 kHz in Run 10 and beyond in both p+p and Au+Au collisions unless vigorous development of needed 10 Gb/s Ethernet technology is begun concurrently with the second generation Data Collection Module (DCM II) now being designed. The new silicon vertex detectors will at least double the event size, and many physics topics of interest are not accessible to a Level 1 trigger. Thus, it is crucial to the ultimate PHENIX physics output to balance a continued large minimum bias recorded bandwidth, while at the same time improving select Level 1 triggers.

1 Introduction

PHENIX Data Acquisition (DAQ) and trigger capabilities have consistently been improved and upgraded over the course of the first eight RHIC runs in order to take advantage of new detectors and improved luminosity. PHENIX data acquisition and triggering capability has kept pace or exceeded the delivered RHIC luminosity both in heavy ion collisions and in polarized p+p. Improvements have come from using event buffering in the Front End Modules (FEM's), code optimization and improved zero suppression in the DCM's, and hardware and software improvements to the Event Builder (EvB) which have kept pace or exceeded increases in RHIC luminosity.

In Au+Au heavy ion collisions, it has been possible up to now to record the vast majority of all minimum bias collisions (via Beam-Beam triggers). This has been enormously successful in providing high statistics for a full suite of observables (many of which do not have Level 1 trigger options, for example, low mass dieleptons and virtual direct photons and many others). Figure 1 shows the livetime fraction as a function of input raw trigger rate for RHIC in Run7 Au+Au running. Typical running at the end of Run 7 was 4.5 kHz. The possibility of achieving $\sim$ 6.5

*Cheng-Yi Chi, Mickey Chiu, Brian Cole, Ed Desmond, John Haggerty, John Lajoie, Alex Linden-Levy, Jamie Nagle, Chris Pinkenburg, Martin Purschke, Carla Vale, David Winter, Matt Wysocki kHz for Au+Au in Run 9 and improvements beyond are detailed below. As we describe in this document, RHIC luminosities will soon exceed this model of running.

In p+p, d+Au, and light ion running, the luminosity significantly exceeds the DAQ bandwidth. However, the low detector occupancy allows the use of a suite of Level 1 triggers. Previous running has included DAQ throughput of order 5-6 kHz in p+p, and we believe that $\sim 7-8$ kHz can be achieved in Run 9 (without upgrade detectors). Significantly new development will be needed to not lose these gains with the large data volume new detectors.

In order to maintain and extend this progress in future RHIC runs, continued development of the DAQ is needed in order to take advantage of increased luminosity due to the application of beam cooling techniques in RHIC, and because a number of PHENIX upgrade detectors are capable of producing a substantial amount of data with new DAQ electronics that will require new hardware and software. Until now, data acquisition development has been done as part of the operating program, but a number of new PHENIX detectors require faster and more complex readout and increased RHIC luminosity coupled with the larger data volume of the new detectors signal the need for some dedicated development and integration with the existing data acquisition infrastructure. In addition, the data acquisition hardware is aging, and the upgrade detectors installed up to now have relied on using spares which are difficult to replace because many components have become hard to find or obsolete.

For these reasons, it is important to begin funding development in the following areas in order of priority:

1. Additional data handling capacity for Run 9 to handle higher Au+Au luminosity.

2. DCM II hardware and software development.

3. EvB hardware upgrades and R&D on 10 Gb/s Ethernet necessary to make effective use of DCM II.

4. R&D on development of improved triggers, including upgrades to the electromagnetic trigger that may necessitate new calorimeter electronics.

5. Tests and development necessary to implement "demultiplexing" all detector electronics which would enable data acquisition rates in the range of 10-20 kHz.

The goal of 1-3 is to be able to acquire data from the present PHENIX detectors and the currently funded upgrade detectors to record data slightly faster than has been achieved in Run 8, about 8 kHz of events in all species. The goal of 4 is enable PHENIX to trigger in future runs when the collision rate exceeds 8 kHz, and the goal of 5 is to determine the feasibility of recording data up to 16 kHz to cope with increases in luminosity and and physics topics without Level 1 trigger capability.

2 Expected rate and data volume

In this section, we will interpret the expectations for RHIC luminosity provided by Fischer et al. in light of the PHENIX upgrade program. Figure 2 shows the expectations for high energy Au+Au collisions.[1] Compared to the Run 7 Au+Au run, the peak luminosity is expected to be 2.4 times

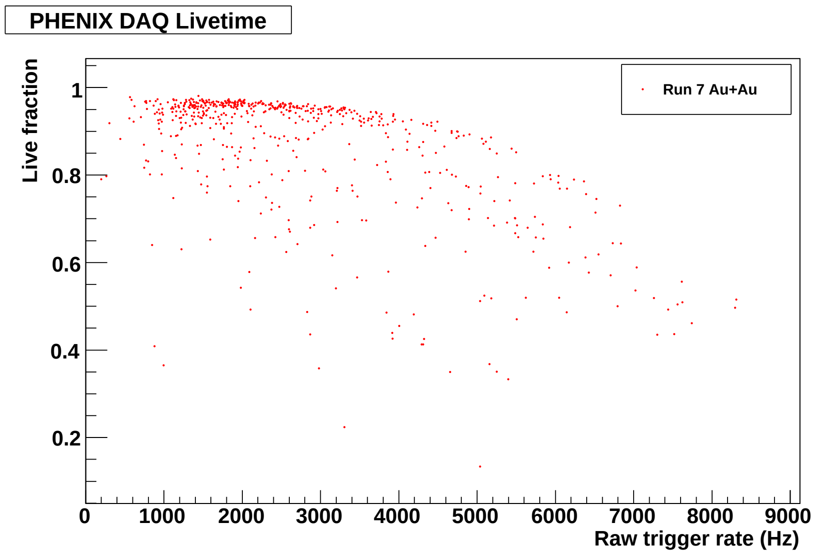

Figure 1: Livetime fraction measured as a function of the minimum bias trigger rate for PHENIX physics runs in Run 7 Au+Au collisions.

as high by Run 13, and the average luminosity in a store is expected to be 4.1 times as large as Run 7 (the larger average is due to increased beam lifetime at store).

From the predicted luminosity, we have projected event rates during a fill as shown in Figure 3. This shows the rate of minimum bias triggers with a vertex within ±30 cm of the center of the PHENIX detector. We assume that the first 30 minutes of the fill are not useful for physics due to beam steering and collimation in RHIC needed to reduce backgrounds to an acceptable level, and recovering from the clock glitches, ramping the high voltage in the wire chambers, and adjusting trigger prescales in PHENIX. As the luminosity and beam lifetime increase, a shrinking fraction of the minimum bias triggers can be recorded as shown in Table 1.

It should be noted that there are any number of possible strategies for sampling a larger fraction of the delivered luminosity, including the "rare" electron and muon triggers, tightening the vertex constraint (which may be desirable for the vertex detectors described in the next section in any case), or selecting more central events, but Table 1 shows that PHENIX will not be able to record basically the entire collision rate beyond Run 9.

With stochastic cooling, the CAD projection is for 86% of collisions with ±30 cm (which is included in the values of Figure 2). However, only 33% of collisions will be within ±10 cm (the region of optimal acceptance of the inner silicon detectors). Thus, for example, for $B \rightarrow J/\psi$ the effective collision rate (within the 10 cm), will be 38% of that shown in Figure 2. However, for J/ψ alone (for J/ψ v_2, for example), the full 30 cm is usable. More detailed studies are needed to determine if the vertex limitation for charm/beauty displaced vertices is a significant problem, given that some of these topics do not currently have Level 1 trigger capability. Note that the electron trigger rejection at low to modest p_T in Au+Au is very small.

Table 4: Delivered RHIC luminosities of the last three Au-Au runs and projected Au-Au luminosities for 100 GeV/nucleon beam energy. Physics runs are assumed to be 12 weeks long in FY 2009 to FY 2013.

Parameter	Unit	FY2002	2004	2007	2009E	2010E	2011E	2012E	2013E
No of bunches	…	55	45	103	103	111	111	111	111
Ions/bunch, initial	10^9	0.6	1.1	1.1	1.1	1.1	1.1	1.20	1.25
Avg. beam current/ring	mA	33	49	112	112	121	121	132	137
β^*	m	1.0	1.0	0.80	0.70	0.60	0.50	0.50	0.50
Hour glass factor	…	0.96	0.96	0.94	0.93	0.91	0.88	0.88	0.88
Beam-beam param./IP	10^{-3}	0.9	1.7	1.5	1.5	1.5	1.5	1.6	1.7
Peak luminosity	10^{26} cm^{-2}s^{-1}	5	15	30	39	48	55	66	72
Avg./peak luminosity	%	30	33	40	43	57	72	70	70
Avg. store luminosity	10^{26} cm^{-2}s^{-1}	1.5	5.0	12	17	27	40	46	50
Time in store	%	25	53	48	60	60	60	60	60
Max. luminosity/week	μb^{-1}	25	160	380	610	980	1,450	1,680	1,820
Min. luminosity/week	μb^{-1}				380	380	380	380	380
Max. luminosity/run	nb^{-1}	0.09	1.37	3.3	6.4	10.3	15.2	17.6	19.1
Min. luminosity/run	nb^{-1}				3.3	3.3	3.3	3.3	3.3

Figure 2: Expectations for Au+Au luminosity.[1]

Year	4.5 kHz	8.0 kHz	15.0 kHz
2007	86%	100%	100%
2009	48%	74%	97%
2010	31%	55%	87%
2011	20%	36%	68%
2012	17%	30%	57%
2013	16%	28%	52%

Table 1: Fraction of minimum bias events that can be recorded for data acquisition rates of 4.5, 8, and 15 kHz for the luminosity profiles shown in Figure 3.

3 New detectors

PHENIX is in the midst of an ambitious detector upgrade program which will bring to PHENIX powerful new capabilities in vertex reconstruction, muon triggering and reconstruction, and increased coverage in electromagnetic calorimetry. These upgrade detectors bring with them increases in data volume that must be recorded and although the upgrade designs have taken account of the need to maintain or increase the data acquisition speed, a significant amount of additional data acquisition hardware and software will be necessary to accomplish this.

In Au+Au collisions in Run 7, the event size for minimum bias triggers was about 200 kbyte, although this ranged from about 90 kbyte for peripheral collisions (or in p+p collisions) up to almost 400 kbyte in the 10% most central events. At 4.5 kHz, that corresponds to about 900 Mbyte/s of recorded data, and compression in the EvB reduces that to about 600 Mbyte/s. That

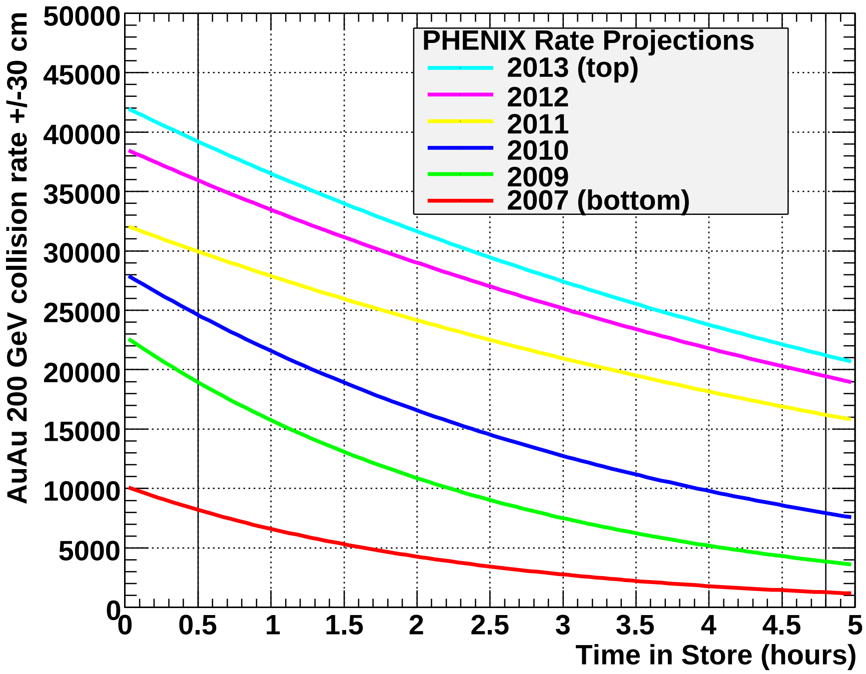

Figure 3: Projected Au+Au collision rate in PHENIX.

speed is the reason that events are recorded locally at PHENIX in six “Buffer Boxes” equipped with Gigabit Ethernet which can record data at that rate. At earlier stages of the data transmission, there are additional speed limitations which can be accommodated by the parallelism of DCM's and EvB, but the new detectors can produce more data than is practical with one gigabit transmission speeds.

The new detectors are capable of producing as much data or more than all the present detectors combined, and the noise rates and zero suppression of the detectors is crucial to operating them at high rates.

VTX Pixels The VTX Pixel detector consists of about 3.9 Mpixel, and although the occupancy even in central Au+Au collisions is expected to be small (0.54% and 0.16% in the inner and outer layers, respectively, according to simulation done for the CDR), the data rate is dependent on the noise rate and the details of the zero suppression scheme. The expectation is that the VTX Pixels will contribute about another 90 kbyte to the event size in p+p or Au+Au collisions (since most of the data comes from irreducible noise hits) when the zero suppression algorithm can be optimized for the noise and background environment. The design calls for 60 DCM II fibers.

VTX Strips The VTX Strip detector consists of about 380,000 strips read out by an 8 bit ADC in the SVX4 ASIC. The occupancy of the strips is expected to be higher than the pixels (4.5% and 2.5% in the inner and outer layers, according to the CDR) and with similar assumptions for the noise and background, the VTX strips should ultimately contribute about 39 kbyte/event. The design calls for 44 DCM II fibers.

FVTX The Forward Vertex detector consists of 4 layers per end of pixel detectors read out with a variant of the FPIX chip developed for the BTeV experiment. Similar assumptions about noise and background lead to an estimate of about 100 kbyte/event from the FVTX in Au+Au collisions. The design calls for 84 DCM II fibers.

NCC The Nosecone Calorimeter (NCC) consist of about 16,000 channels of 14 bit ADC in two arms. The occupancy in Au+Au collisions is likely to be quite high, although the expectations are quite preliminary, we will assume 50% occupancy, which would contribute 32 kbyte/event. The design calls for 60 DCM II fibers.

Thus, the presently envisioned new detectors are expected to produce about 200-300 kbyte/event, and even the most optimistic assumptions would lead us to expect that the total event size will be doubled.

4 DCM II and the event builder

The readout of the DCM II has not yet been developed, but the design presently envisioned would be through a custom PCI Express interface read into a conventional Intel or AMD motherboard (called a Sub-Event Buffer, or SEB) that can package the data and transmit it into the EvB over a 10 Gb/s Ethernet network. The interface to the DCM II will be designed so that there is no bottleneck at the SEB, that is, if data can be written to the 10 Gb/s Ethernet network at 10 Gb/s, data can be received from the DCM II's at 10 Gb/s as well. (The DCM's can be read into the present generation of SEB at about 80% of 1 Gb/s, and transmit data onto the EvB network at about the same speed.) Table 2 shows a possible segmentation of the DCM II's into SEB's and the data rates from the typical zero suppressed event sizes described above. The actual arrangement will probably be different, based on the actual event sizes and other operational considerations, like the need for expeditious pedestal runs, but this scenario gives a reasonable estimate for the number of machines that must be accommodated, and shows the need for speeds beyond 1 Gb/s for these detectors.

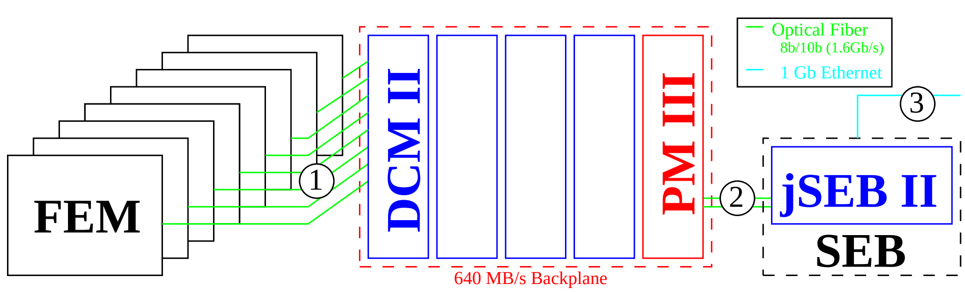

Figure 4: Block diagram of conceptual design of DCM II and their interface to the Event Builder.

System	Groups	Event size (kbyte)	Data rate (Gb/s)
VTX Pixel	3	90	1.92
VTX Strip	2	39	1.25
FVTX	6	100	1.07
NCC	2	32	1.02

Table 2: Groups of DCM's needed for upgrade detectors with the expected event size from each and the data rate at 8 kHz.

5 Run 9

Although this proposal is aimed at longer term improvements to data acquisition and triggering, it should be noted that a number of presently anticipated improvements to data acquisition could require data logging capacity even for Run 9, and other changes and improvements. For example, it may be necessary to add additional SEB's to handle data from the central drift chamber, which currently is close to saturating the available bandwidth, and the buffer boxes almost certainly must be augmented or reconfigured to record data faster than 500 Mbyte/s.

6 Trigger upgrades

Level 1 triggers are crucial to future PHENIX running. Although much of the PHENIX physics program depends on recording a large sample of events for which there is no effective trigger, electron and muon triggers are crucial rare event triggers.

An upgrade to the muon trigger is being developed and deployed in PHENIX in the near future and is expected to become operational in Run 9. Since that project is underway and has been separately proposed and reviewed, we will not discuss muon triggers further, although a number of smaller development projects may be necessary to fully realize muon triggering beyond what is needed for triggering on muons from $W^{\pm}$ decays.

PHENIX has effectively utilized Level 1 trigger for muons in p+p, d+Au, and light ion running. The current trigger configuration provides insufficient rejection for single very high p_T muons (from W and Z boson decays). This crucial part of the spin program is being enabled by the forward muon trigger upgrades. These upgrades are already funded and underway and are thus not described further here. Additional modest muon trigger upgrades to allow panel-based muon tracking (possibly implemented in the existing hardware with new firmware) is being investigated. This and other noise reduction are important for heavy ion triggering in the future.

The PHENIX central arm electromagnetic trigger could be improved. The present trigger is subject to noise and calibration problems of the electromagnetic calorimeters which limit the trigger rejection that can be achieved with it. Although this does not limit the PHENIX physics program at this time, in future runs where less of the total cross section will be recorded, an electron trigger with fewer dead regions and a sharper turn-on in energy and the ability to set a lower threshold on the 4×4 overlapping tile trigger will be needed.

Triggering in the electromagnetic calorimeter should be studied and development of improvements to the trigger should begin soon. It is worth considering the larger project of developing new calorimeter electronics which could be designed with a fully digital Level 1 trigger using a pipelined converter ADC which could use digital data for triggering before it enters the data stream. The HBD electronics provide a possible prototype for this approach, and a relatively small amount of development and testing would suggest whether to pursue this approach to a proposal for new calorimeter electronics. A block diagram of the realization of such a design is shown in Figure 5.

The VTX Pixel and VTX Strip detectors will cover about ±10 cm around the vertex, and in order to operate them efficiently at higher luminosities, it will be necessary to modify the Beam-Beam Local Level 1 (BBCLL1) trigger so that multiple vertex cuts can be made. Several approaches to this project are possible, largely using the existing hardware, but since the BBCLL1 trigger is so central to the PHENIX physics program, it is important to develop and debug this trigger carefully but expeditiously.

It should be noted that the NCC project does not include a Level 1 trigger, so that there is presently no calorimeter trigger that covers the NCC rapidity region, although the NCC electronics is designed with the possibility of a Level 1 trigger. It is important to determine soon whether a Level 1 trigger must be developed for the NCC.

Last, it should be noted that all PHENIX triggers are orchestrated by the "Global Level 1" (GL1) trigger. This system consists of aging VME boards, and limit PHENIX to a total of 32 triggers. The aging hardware, coupled with the limitation on the total number of triggers and other limitations suggest that it may be necessary begin planning for maintenance, modification, or replacement of parts of this system.

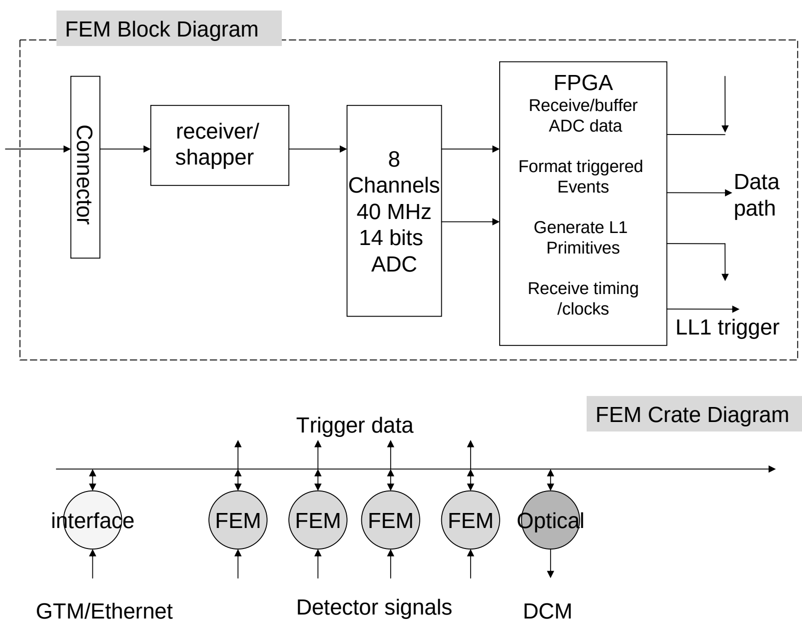

Figure 5: Block diagram of possible new EMCAL electronics.

7 Demultiplexing

The PHENIX detector electronics was designed to allow readout of any Front End Module in 40 μs by the DCM's, but in order to reduce the costs of fiber transmitters, optical fibers, and DCM's, two FEM's are multiplexed to a single DCM channel in most PHENIX systems. Ultimately, this is the step that limits the highest speed that PHENIX detectors can be read. With the FEM's multiplexed two and the ADC conversion times optimally overlapped with the readout, that limits the maximum Level 1 trigger rate to 11.9 kHz, which is the highest rate that could ever be achieved with multiplexed FEM's.

Provision was made for demultiplexing the FEM's by adding fiber transmitters and DCM channels in the design of the PHENIX electronics. It would be difficult to demultiplex the entire PHENIX detector, both because of cost and access to the FEM's, particularly the Muon Tracker FEM's, which have over 350 FEM's installed inside the north and south muon magnets, but it may be possible to modify the PHENIX trigger so that the central arm detectors could be demultiplexed and the muon arms are not, allowing higher rates in central arm triggers than in the muon arms. Even so, not all central arm detectors are capable of conversion within 40 μs, and our present understanding of the conversion time of the electromagnetic calorimeter electronics would limit the maximum possible central arm readout to 17.2 kHz.

Although we do not have a plan to implement demultiplexed readout in the foreseeable future, it is potentially important to begin tests, measurements, and possibly some development for how to implement demultiplexed detectors in the future. This would involved tests and possible modifications to the Global Level 1 (GL1) trigger, and testing of all FEM systems to determine whether they perform adequately without multiplexing, and possibly development of a new generation of optical daughterboards which would enable readout of old detectors by DCM II's. Without starting on these longer term developments, it will surely not be possible to readout PHENIX above 10 kHz.

8 Future work

As discussed in the introduction and explained above, a number of lines of development must be pursued in order to accommodate new detectors and maximize the readout speed of the entire detector in the near and more distant future. Implementing the readout of the new detectors will be chaotic and interfere with the operation of the PHENIX detector unless we begin to develop hardware and software in a laboratory environment before attempting to deploy it in PHENIX.

1. Additional data handling capacity for Run 9 to handle higher Au+Au luminosity.
2. DCM II hardware and software development.
   (a) Development of the prototype DCM II as rapidly as possible
   (b) Purchase development licenses for driver tools as necessary
   (c) Purchase test systems with available PCI Express slots and 10 Gb/s Ethernet cards
3. EvB hardware upgrades and R&D on 10 Gb/s Ethernet necessary to make effective use of DCM II.
   (a) Purchase test systems with 10 Gb/s Ethernet

   (b) Purchase 10 Gb/s Ethernet capable switching equipment

4. R&D on development of improved triggers, including upgrades to the electromagnetic trigger that may necessitate new calorimeter electronics.

   (a) Adapt HBD or NCC electronics to read out one or two EMCAL supermodules
   (b) Develop conceptual design of an improved electromagnetic calorimeter trigger

5. Tests and development necessary to implement "demultiplexing" all detector electronics which would enable data acquisition rates in the range of 10-20 kHz.

   (a) Test existing FEM systems for demultiplexed readout
   (b) Develop GL1 capability to run two partitions with different rates sharing some detectors
   (c) Develop optical daughterboard compatible with DCM II

9 Requests for R&D resources

In order to prepare a full proposal for a substantial revision to the PHENIX DAQ and trigger, the DAQ and trigger groups will need to assemble prototypes and carry out tests with the prototype DCM II's as they become available. This development environment can be assembled so that it will provide a useful test bench for hardware and software development and debugging, and during the deployment of the new detectors, it will provide a test setup for bringing the new detectors into operation, and ultimately it will provide a debugging facility for the installed detectors. This test facility can also be extended to support repair of existing electronics.

For trigger development, a number of simulation tasks must be pursued, but design of a BBCLL1 capable of taking advantage of the vertex detectors should be begun with high priority, because the usefulness of the new vertex detectors would be severely compromised without this.

For trigger development, we specifically recommend the following

Vertex cuts Begin design of an improved BBCLL1 trigger capable of multiple vertex cuts which will enable selection of events within the range accessible to the new vertex detectors (± 10 cm) in addition to continuing to collect (with a different prescale factor) minimum bias triggers with the present vertex cuts ($\sim \pm 30$ cm). Initial design studies are needed soon, which will then necessitate engineering expenditures of order \$50k to ISU and Ames Laboratory.

MUID trigger improvements It is likely that firmware improvements to the MUIDLL1 can produce additional rejection and immunity to noise. Design studies should be done which may result in a request for engineering support at ISU and Ames.

Electron trigger improvements We recommend a several pronged attempt to improve the electron trigger:

- Monte Carlo studies to clarify how much rejection could be reasonably expected from the present EMCAL-RICH trigger with the existing detectors and improved electronics.
- Study of improvements to the existing EMCAL and RICH hardware.

- Design of prototype electronics for the EMCAL which could improve its trigger capability; we estimate that about \$25k of engineering support at Columbia University is needed for this.

For DAQ development, we recommend that a small test cluster be constructed in a laboratory environment separate from the operating PHENIX DAQ in order to avoid disruption of PHENIX datataking.

DCM II Development of the DCM II requires \$380k (Columbia) and \$40k (Colorado) in FY2008 in order for its development to keep pace with the detectors its serves.

Network equipment A small 10 Gb/s Ethernet switch should be purchased for development and testing. An eight port Cisco or Foundry switch costing estimated at \$20k should be adequate, although the available switches should be compared before purchase.

SEB's and ATP's One or two machines which can function as prototype SEB's and ATP's should be purchased, along with 10 Gb/s Ethernet cards. (ATP's are not expected to require 10 Gb/s connectivity, but a machine that can serve as an ATP in testing is needed, and it can serve as an SEB in the future.) The cost is estimated to be \$20k.

Buffer box A machine to be used as a buffer box should be purchased to benchmark logging performance on a 10 Gb/s Ethernet network. Although a large amount of disk space is not needed for testing, it may be possible to move this machine into operation ultimately, and so purchasing a machine and disk array with flexibility and expandability is desirable. Such a machine is estimated to cost \$20k.

DCM daughterboards Development of an improved daughterboard for the DCM should be investigated. An engineering prototype would probably cost about \$20k to develop by reusing some existing parts, after investigating the feasibility of making substantial improvements with such a design.

References

[1] W. Fischer, T. Roser, M. Bai, K. Brown, H. Huang, and C. Montag. *RHIC Collider Projections (FY2009-FY2013)*, February 2008.